



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/894,310	06/27/2001	Richard L. Coulson	42390P11448	7622

8791 7590 11/03/2003

BLAKELY SOKOLOFF TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD, SEVENTH FLOOR
LOS ANGELES, CA 90025

EXAMINER

INOA, MIDYS

ART UNIT	PAPER NUMBER
----------	--------------

2188

DATE MAILED: 11/03/2003

10

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

09/894,310

Applicant(s)

COULSON ET AL.

Examiner

Midys Inoa

Art Unit

2188

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 28 July 2003.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-36 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-36 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 27 July 2001 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
- 11) ☐ The proposed drawing correction filed on _____ is: a) ☐ approved b) ☐ disapproved by the Examiner.
If approved, corrected drawings are required in reply to this Office action.
- 12) ☐ The oath or declaration is objected to by the Examiner.

Priority under 35 U.S.C. §§ 119 and 120

- 13) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
* See the attached detailed Office action for a list of the certified copies not received.
- 14) ☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. § 119(e) (to a provisional application).
a) ☐ The translation of the foreign language provisional application has been received.
- 15) ☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. §§ 120 and/or 121.

Attachment(s)

- 1) ☐ Notice of References Cited (PTO-892) 4) ☒ Interview Summary (PTO-413) Paper No(s). 10.
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948) 5) ☐ Notice of Informal Patent Application (PTO-152)
- 3) ☐ Information Disclosure Statement(s) (PTO-1449) Paper No(s) _____ 6) ☐ Other: _____

DETAILED ACTION

Claim Rejections - 35 USC § 103

1. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

2. Claims 1, 2, 4, 7, 10, 11, 13, 14, 30, 33, and 34 are rejected under 35 U.S.C. 103(a) as being unpatentable over Maheshwari (5,974,508) in view of Teoman et al. (6,463,509).

Regarding Claims 1, 2, 30, 33, and 34, Maheshwari teaches a CPU with a cache memory system in which data is stored in the cache when a cache miss occurs and data must be staged from main memory ("storing data in first memory"). This cache system has the ability to enter a local locking mode ("pinning"), which allows the cache to lock selected time-critical entries in the cache to prevent them from being replaced when new data is being staged from main memory (Column 4, lines 33-42). Maheshwari does not teach the use of a non-volatile cache. Teoman et al. teaches a non-volatile cache large enough to hold several hundred megabytes worth of data ("mass storage cache", Column 3, lines 4-6). It would have been obvious by one of ordinary skill in the art at the time the invention was made to replace the cache of Maheshwari with the non-volatile mass storage cache of Teoman et al. because a non-volatile mass storage cache has the advantage of not losing data during power loss and it can hold larger amounts of data than a regular cache, thus allowing it to be used as a boot source (Teoman et al., Column 3, lines 6-10). Although replacing the existing cache of Maheshwari for a slower non-volatile cache might slow down the system, Maheshwari's intended method of processing time reduction

Art Unit: 2188

would still reduce running time (when compared to another system with a non-volatile cache) and the advantages of using a non-volatile cache as opposed to a regular cache is a tradeoff valuable enough to be worth the small slow down being caused in the system.

Regarding Claims 4 and 7, Maheshwari teaches the storing of a least recently used bit (LRU) and an entry lock bit ("metadata... corresponding line of data is pinned") in a tag status field of the cache memory (Column 4, line 62 –Column 5, lines 7).

Regarding Claims 10-11 and 13, Maheshwari teaches a CPU with a cache memory system in which data is stored in the cache when a cache miss occurs and data must be staged from main memory ("storing data in first memory"). This cache system has the ability to enter a local locking mode ("pinning"), which allows the cache to lock selected time-critical entries in the cache to prevent them from being replaced when new data is being staged from main memory (Column 4, lines 33-42). Maheshwari also teaches the storing of a least recently used bit (LRU) and an entry lock bit ("metadata... corresponding line of data is pinned") in a tag status field ("second storage medium included in the cache") of the cache memory (Column 4, line 62 –Column 5, lines 7). Maheshwari does not teach the use of a non-volatile cache.

Teoman et al. teaches a non-volatile cache large enough to hold several hundred megabytes worth of data ("mass storage cache", Column-3, lines 4-6). It would have been obvious by one of ordinary skill in the art at the time the invention was made to replace the cache of Maheshwari with the non-volatile mass storage cache of Teoman et al. because a non-volatile mass storage cache has the advantage of not losing data during power loss and it can hold larger amounts of data than a regular cache, thus allowing it to be used as a boot source (Teoman et al., Column 3, lines 6-10). Although replacing the existing cache of Maheshwari for a slower non-volatile

Art Unit: 2188

cache might slow down the system, Maheshwari's intended method of processing time reduction would still reduce running time (when compared to another system with a non-volatile cache) and the advantages of using a non-volatile cache as opposed to a regular cache is a tradeoff valuable enough to be worth the small slow down being caused in the system.

Regarding Claim 14, the combination of Maheshwari in view of Teoman et al. does not teach the use of an add-in card as a cache. It would have been obvious at the time the invention was made to use a flash-card ("add-in card") to replace the cache of the invention in the combination of Maheshwari in view of Teoman et al. because flash cards are a type of non-volatile memory and a device of this type would allow the addition of a supplemental portable non-volatile cache; which is very convenient in portable systems.

3. Claims 3, 8, 31, 32, and 35 are rejected under 35 U.S.C. 103(a) as being unpatentable over Maheshwari (5,974,508) in view of Teoman et al. (6,463,509) as applied to claims 1, 7, 30, and 33 above, and further in view of Berstis (US 2002/0174370 A1).

Regarding Claims 3, 31, and 35, the combination of Maheshwari in view of Teoman et al. does not teach locking data needed for system initialization. Berstis teaches locking program code and data needed for power-up notifications (Page 4, paragraph 52). It would have been obvious to one of ordinary skill in the art at the time the invention was made to integrate Berstis' locking of power-up data with the invention of the combination of Maheshwari in view of Teoman et al. because it would allow this data to be available faster for the power up procedure. In addition, Maheshwari teaches locking of time-critical entries, it is understood that data needed for power up constitutes as a time-critical entry.

Art Unit: 2188

Regarding Claim 32, it is understood that during a power-up procedure, a system's resources are focused on the initialization of the system and thus, pinning of other data needed for other purposes would be limited to a predetermined maximum amount of data that can be pinned in order to prevent resources from being used in unnecessary processes.

Regarding Claim 8, the combination of Maheshwari in view of Teoman et al. does not teach pinning data needed for system initialization or a third state in the metadata to indicate if a line of data was present in the cache before system initialization. Berstis teaches locking program code and data needed for power-up notifications (Page 4, paragraph 52). It would have been obvious to one of ordinary skill in the art at the time the invention was made to integrate Berstis' locking of power-up data with the invention of the combination of Maheshwari in view of Teoman et al. because it would allow this data to be available faster for the power up procedure. In addition, it would have been obvious to create a third state in the metadata to categorize lines of data present before the system initialization since such a status bit would allow the system to better identify data that is necessary during the power-up process.

4. Claims 5, 6, 9, 12, and 36 are rejected under 35 U.S.C. 103(a) as being unpatentable over Maheshwari (5,974,508) in view of Teoman et al. (6,463,509) as applied to claims 1, 4, 7, and 35 above, and further in view of Shirata et al. (US 2001/0043784 A1).

The combination of Maheshwari in view of Teoman et al. does not teach storing metadata in a second memory. Shirata et al. teaches a first and second memory which need a power supply or a backup power source to keep stored contents ("volatile"), in which metadata is stored separately in the second memory (Page 2, Paragraph 0028). It would have been obvious to one of ordinary skill in the art at the time the invention was made to integrate the separate storing of

metadata in Shirata et al. with the system in the combination of Maheshwari in view of Teoman et al. because storing the metadata separately allows the system to sort out the data by type (metadata vs. cached data) thus making the storage system more organized.

5. Claims 15-29 are rejected under 35 U.S.C. 103(a) as being unpatentable over Maheshwari (5,974,508) in view of Teoman et al. (6,463,509) and further in view of Berstis (US 2002/0174370 A1).

Regarding Claims 15-16, Maheshwari teaches a CPU with a cache memory system in which data is stored in the cache when a cache miss occurs and data must be staged from main memory ("storing data in first memory"). This cache system has the ability to enter a local locking mode ("pinning"), which allows the cache to lock selected time-critical entries in the cache to prevent them from being replaced when new data is being staged from main memory (Column 4, lines 33-42). Maheshwari does not teach the use of a non-volatile cache or the locking data needed for system initialization. Teoman et al. teaches a non-volatile cache large enough to hold several hundred megabytes worth of data ("mass storage cache", Column 3, lines 4-6). Berstis teaches locking program code and data needed for power-up notifications (Page 4, paragraph 52). It would have been obvious by one of ordinary skill in the art at the time the invention was made to replace the cache of Maheshwari with the non-volatile mass storage cache of Teoman et al. because a non-volatile mass storage cache has the advantage of not losing data during power loss and it can hold larger amounts of data than a regular cache, thus allowing it to be used as a boot source (Teoman et al., Column 3, lines 6-10). Although replacing the existing cache of Maheshwari for a slower non-volatile cache might slow down the system, Maheshwari's intended method of processing time reduction would still reduce running time

Art Unit: 2188

(when compared to another system with a non-volatile cache) and the advantages of using a non-volatile cache as opposed to a regular cache is a tradeoff valuable enough to be worth the small slow down being caused in the system. Additionally, it would have been obvious to one of ordinary skill in the art at the time the invention was made to integrate Berstis' locking of power-up data with the invention of the combination of Maheshwari in view of Teoman et al. because it would allow this data to be available faster for the power up procedure. Since Maheshwari teaches locking of time-critical entries, it is understood that data needed for power up constitutes as a time-critical entry.

Regarding Claim 17, it is understood that during a power-up procedure, a system's resources are focused on the initialization of the system and thus, pinning of other data needed for other purposes would be limited in order to prevent resources from being used in unnecessary processes.

Regarding Claim 18, Maheshwari teaches the storing of a least recently used bit (LRU) and an entry lock bit ("metadata... corresponding line of data is pinned") in a tag status field of the cache memory (Column 4, line 62 –Column 5, lines 7).

Regarding Claim 19, the combination of Maheshwari in view of Teoman et al. further in view of Berstis does not teach setting a timer upon system initialization and setting a first state indicating that the data is pinned only until the time expires. It would have been obvious to one of ordinary skill in the art at the time the invention was made to set a timer in order to prevent the system initialization data from being pinned longer than it is needed.

Regarding Claims 20 and 26, it is understood that during a power-up procedure, a system's resources are focused on the initialization of the system and thus, pinning of other data

needed for other purposes would be limited to a predetermined maximum amount of data that can be pinned in order to prevent resources from being used in unnecessary processes.

Regarding Claim 21, the combination of Maheshwari in view of Teoman et al. further in view of Berstis does not teach setting a second state for data that was present before system initialization or setting a timer upon system initialization and setting a first state indicating that the data is pinned only until the time expires. It would have been obvious to one of ordinary skill in the art at the time the invention was made to create a second state in the metadata to categorize lines of data present before the system initialization since such a status bit would allow the system to better identify data that is necessary during the power-up process and discard data that is not needed. In addition, it would have been obvious to one of ordinary skill in the art at the time the invention was made to set a timer in order to prevent the system initialization data from being pinned longer than it is needed. It is understood that during a power-up procedure, a system's resources are focused on the initialization of the system and thus, pinning of other data needed for other purposes would be limited to a predetermined maximum amount of data that can be pinned in order to prevent resources from being used in unnecessary processes.

Therefore, in these circumstances, a system should only pin data when the maximum amount of data has not been reached and the timer set has not expired. In addition, when the maximum amount of data has been reached and the timer has expired, if the second state of a line of data is not set, then this data has be cleared since it has been identified as not needed for system initialization and at this time, data needed for power-up takes priority.

Regarding Claim 22, Maheshwari teaches the storing of a least recently used bit ("age of a line of data") in a tag status field of the cache memory (Column 4, line 62 –Column 5, lines 7).

Art Unit: 2188

The LRU bit is known to be used as a criteria discarding data that when space is needed in a cache.

Regarding Claims 23, 25 and 28, Maheshwari teaches a CPU with a cache memory system in which data is stored in the cache when a cache miss occurs and data must be staged from main memory ("storing data in first memory"). This cache system has the ability to enter a local locking mode ("pinning"), which allows the cache to lock selected time-critical entries in the cache to prevent them from being replaced when new data is being staged from main memory (Column 4, lines 33-42). Maheshwari teaches the storing of a least recently used bit (LRU) and an entry lock bit ("metadata... corresponding line of data is pinned") in a tag status field of the cache memory ("second storage media... included in the cache", Column 4, line 62 –Column 5, lines 7). Maheshwari does not teach the use of a non-volatile cache or the locking data needed for system initialization. Teoman et al. teaches a non-volatile cache large enough to hold several hundred megabytes worth of data ("mass storage cache", Column 3, lines 4-6). Berstis teaches locking program code and data needed for power-up notifications (Page 4, paragraph 52). It would have been obvious by one of ordinary skill in the art at the time the invention was made to replace the cache of Maheshwari with the non-volatile mass storage cache of Teoman et al. because a non-volatile mass storage cache has the advantage of not losing data during power loss and it can hold larger amounts of data than a regular cache, thus allowing it to be used as a boot source (Teoman et al., Column 3, lines 6-10). Although replacing the existing cache of Maheshwari for a slower non-volatile cache might slow down the system, Maheshwari's intended method of processing time reduction would still reduce running time (when compared to another system with a non-volatile cache) and the advantages of using a non-

Art Unit: 2188

volatile cache as opposed to a regular cache is a tradeoff valuable enough to be worth the small slow down being caused in the system. Additionally, it would have been obvious to one of ordinary skill in the art at the time the invention was made to integrate Berstis' locking of power-up data with the invention of the combination of Maheshwari in view of Teoman et al. because it would allow this data to be available faster for the power up procedure. Since Maheshwari teaches locking of time-critical entries, it is understood that data needed for power up constitutes as a time-critical entry.

Regarding Claim 24, the combination of Maheshwari in view of Teoman et al. further in view of Berstis does not teach setting a second state for data that was present before system initialization. It would have been obvious to create a third state in the metadata to categorize lines of data present before the system initialization since such a status bit would allow the system to better identify data that is necessary during the power-up process.

Regarding Claim 27, the combination of Maheshwari in view of Teoman et al. further in view of Berstis does not teach storing metadata in a second volatile memory. Shirata et al. teaches a first and second memory which need a power supply or a backup power source to keep stored contents ("volatile"), in which metadata is stored separately in the second memory (Page 2, Paragraph 0028). It would have been obvious to one of ordinary skill in the art at the time the invention was made to integrate the separate storing of metadata in Shirata et al. with the system in the combination of Maheshwari in view of Teoman et al. because storing the metadata separately allows the system to sort out the data by type (metadata vs. cached data) thus making the storage system more organized.

Regarding Claim 29, the combination of Maheshwari in view of Teoman et al. further in view of Berstis does not teach the use of an add-in card as a cache. It would have been obvious at the time the invention was made to use a flash-card ("add-in card") to replace the cache of the invention in the combination the of Maheshwari in view of Teoman et al. because flash cards are a type of non-volatile memory and a device of this type would allow the addition of a supplemental portable non-volatile cache; which is very convenient in portable systems.

Response to Arguments

6. Applicant's arguments filed on July 28th, 2003 regarding independent claims 1, 7, 10, 15, 23, and 30 have been fully considered but they are not persuasive.

Applicant argues that changing the cache of Maheshwari for the non-volatile cache of Teoman would render the invention of Maheshwari unsatisfactory for its intended purpose since a non-volatile cache would introduce delays in a system whose purpose is to reduce memory access time. However, replacing the existing cache of Maheshwari for a slower non-volatile cache would not affect the method being applied to increase performance. Although the performance would not be as high as that of a system with a volatile memory, the method of Maheshwari would still reduce running time in a system with a non-volatile when this system is compared to other systems with non-volatile caches. Additionally, the advantages of using a non-volatile cache as opposed to a regular cache can be considered a valuable tradeoff important enough to be worth the small slow down being endured by the system due to the change in cache.

Art Unit: 2188

Conclusion

7. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire **THREE MONTHS** from the mailing date of this action. In the event a first reply is filed within **TWO MONTHS** of the mailing date of this final action and the advisory action is not mailed until after the end of the **THREE-MONTH** shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than **SIX MONTHS** from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Midys Inoa whose telephone number is (703) 305-7850. The examiner can normally be reached on M-F 7:00am - 4:30pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mano Padmanabhan can be reached on (703) 306-2903. The fax phone number for the organization where this application or proceeding is assigned is (703) 872-9306.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.

Midys Inoa
Midys Inoa
Examiner
Art Unit 2188

Mano Padmanabhan
10/31/03

MANO PADMANABHAN
SUPERVISORY PATENT EXAMINER
TC 2100

MI